arXiv:1611.06739v1 [stat.ME] 21 Nov 2016

Simultaneous Control of All False Discovery Proportions by an Extension of Hommel's Method

Rosa J. Meijer[1,2] Thijmen J.P. Krebs[3] Aldo Solari[4]
Jelle J. Goeman[1,5*]

November 22, 2016

Abstract

We revisit closed testing with local tests based on the Simes inequality that underlies Hochberg's and Hommel's famous methods of familywise error rate (FWER) control. We show that free additional inferences on the false discovery proportion of any subset of the hypotheses may be obtained from this closed testing procedure, where the subset of interest may be selected post hoc. The resulting method will be shown to be related to the method for false discovery rate control of Benjamini and Hochberg. Like that method, and unlike Hommel's original method, the new method does not suffer from diminishing power when the number of hypotheses increases. Computation time also scales well. With m hypotheses it is linear in subset size after an initial $O(m \log(m))$ preparatory step. We also present a new algorithm that calculates adjusted p-values for Hommel's method in $O(m \log(m))$ time, improving upon the current quadratic time algorithm.

1 Introduction

The method of Hommel (1988) is a well-known multiple testing procedure that is uniformly more powerful than the methods of Bonferroni (Dunn, 1961), Holm (1979) and Hochberg (1988). It controls the familywise error rate (FWER), guaranteeing that

[1]Leiden University Medical Center, Department of Medical Statistics and Bioinformatics, Leiden, The Netherlands
[2]HOVON Data Center, Department of Hematology, Erasmus University Medical Center, Rotterdam, The Netherlands
[3]Delft University of Technology, Delft, The Netherlands
[4]University of Milano-Bicocca, Department of Economics, Management and Statistics, Milan, Italy
[5]Radboud University Medical Center, Radboud Institute for Molecular Life Science, Nijmegen, The Netherlands
*Corresponding author. Email: j.j.goeman@lumc.nl

with probability at least $1 - \alpha$ no true null hypotheses are rejected. Hommel's method can be constructed using a combination of the closed testing procedure (Marcus et al., 1976) with local tests based on the inequality of Simes (1986). It is valid whenever the Simes inequality can be assumed to hold for all p-values corresponding to true null hypotheses. This condition is the same as is required for the validity of the False Discovery Rate (FDR) controlling procedure of Benjamini and Hochberg (1995), and has been extensively studied elsewhere (e.g. Benjamini and Yekutieli, 2001; Rødland, 2006; Sarkar, 2008; Finner et al., 2014).

Recently, Goeman and Solari (2011) showed that the closed testing procedure and methods based on it can be used to construct confidence intervals for the number, or equivalently the proportion, of true (or false) hypotheses in any subset of the hypotheses. These confidence intervals are simultaneous over all subsets by construction. Therefore, they allow confidence statements for the false discovery proportion (FDP) of any subset of hypotheses selected post hoc by the user. This in turn, allows the roles of the user and the procedure in multiple testing to be reversed: rather than letting the user choose the error rate before seeing the data and letting the procedure determine the rejected set, the user may select the hypotheses he or she believes to be false after seeing the data. The procedure then gives a confidence statement for the number or proportion of false positives in that set.

In this paper we combine the confidence interval approach of Goeman and Solari (2011) with Hommel's powerful idea of using Simes local tests in the closed testing procedure. The resulting novel multiple testing method will be shown to have a close link to the method of Benjamini and Hochberg (1995), valid under the same assumptions. Perhaps surprisingly, we show that like the method of Benjamini and Hochberg, but unlike Hommel's original method, the new method scales to larger multiple testing problems without loss of power. As the number of hypotheses m increases, the number of rejections allowed at any desired FDP level grows linearly with m, provided a minimal amount of signal is present.

The resulting method is very flexible and can be used in different ways. We distinguish four main types of use. (1.) Control of tail probabilities of the FDP (Genovese and Wasserman, 2006) without having to decide the level of the FDP in advance; (2.) Control of tail probabilities for the FDP when findings are not prioritized by p-values; (3.) Control of tail probabilities of the FDP when exchangeability of hypotheses cannot be assumed; and (4.) All-resolution inference, i.e. simultaneous inference on the amount of signal in all subsets of all sizes of the data (Pacifico et al., 2007).

The combination of closed testing and Simes presents a huge algorithmic challenge for large problems, since naive calculation of the confidence intervals from the closed testing procedure takes exponential time in the number of hypotheses m. Goeman and Solari (2011, 2013) already provided an approximate computational shortcut, related to Hochberg's method, that resulted in a conservative variant of the

method of this paper. We improve upon their method with a fast and exact shortcut. From novel results on the structure of the rejections in the closed testing procedure we calculate confidence statements for the FDP of any set in a time linear in set size, after a common initial preparatory step that takes only $O(m \log(m))$ time. These confidence statements are exact shortcuts, i.e. they are identical to those that would be obtained from the full closed testing procedure.

Additionally, we improve upon the original quadratic time algorithm for calculating FWER-adjusted p-values of Hommel's method (Hommel, 1988), calculating Hommel's adjusted p-values in $O(m \log(m))$ time instead. This new algorithm forces a reassessment of the relative merits of Hommel's method relative to the method of Hochberg (1988). Hochberg's method is a conservative approximation to Hommel's method (Hommel, 1989) that sacrifices power for simplicity and computational efficiency. However, applying Hochberg's method on the sorted p-values also takes $\Theta(m \log(m))$ time.

The structure of the paper is as follows. We start by briefly summarizing the work of Hommel (1988) and Goeman and Solari (2011), constructing a novel faster algorithm for Hommel-adjusted p-values in the process. Next, we derive a simple expression to determine whether an intersection hypothesis will be rejected by the closed testing procedure, from which we derive an expression for the FDP confidence bounds and a fast algorithm to calculate such bounds for a sequence of sets of interest. We show that all interesting results from the procedure are concentrated in a subset of hypotheses related to the rejection set of the procedure of Benjamini and Hochberg (1995), and establish further links with that method. Finally, we investigate the behavior in terms of power of the closed testing procedure as the number of hypotheses $m \to \infty$, showing that unlike in typical FWER-controlling methods, interesting statements do not vanish relative to the size of the multiple testing problem.

2 Closed testing and FDP confidence

We briefly revisit Goeman and Solari (2011) to set the stage and introduce notation.

Suppose we have m hypotheses $H_1, \ldots, H_m$ that we are interested in rejecting. Some of the hypotheses are true, while others are false. Denote by $T \subseteq \{1, \ldots, m\}$ the index set of true hypotheses. The aim of a multiple testing procedure is to come up with an index set $S \subseteq \{1, \ldots, m\}$ of rejected or 'selected' hypotheses, the *discoveries*. This selection should avoid *false discoveries*, which are hypotheses with indices in $T \cap S$. We denote the number of false discoveries for a selected set S as

$$\tau(S) = |T \cap S|,$$

where $|\cdot|$ denotes the size of a set.

Goeman and Solari (2011) showed that it is possible to calculate confidence bounds

$t_\alpha(S)$ such that

$$\mathrm{P}(\tau(S) \leq t_\alpha(S)) \geq 1 - \alpha \tag{1}$$

simultaneously over all $S \subseteq \{1, \ldots, m\}$. Using these confidence bounds a data analyst can study the data, choose any set S of interest and use the multiple testing procedure to obtain a confidence bound for the number of false discoveries in that set. The choice of S can be revised as many times as desired, since the simultaneity of the confidence bounds guarantees their post hoc validity. The bounds can also be formulated in terms of the false discovery proportion (FDP), defined as $\tau(S)/|S|$ if $S \neq \emptyset$, and 0 otherwise. We find

$$q_\alpha(S) = \frac{t_\alpha(S)}{|S|},$$

also defined as 0 if $S = \emptyset$, as an immediate simultaneous upper $(1 - \alpha)$-confidence bound for the FDP. We denote $d_\alpha(S) = |S| - t_\alpha(S)$ as the simultaneous lower $(1-\alpha)$-confidence bound for the number of false null hypotheses $|S| - \tau(S)$.

Bounds $t_\alpha(S)$ can be derived from the closed testing procedure of Marcus et al. (1976). In the general closed testing procedure, the collection of hypotheses is augmented with all possible intersection hypotheses $H_I = \bigcap_{i \in I} H_i$, with $I \subseteq \{1, \ldots, m\}$. An intersection hypothesis H_I is true if and only if H_i is true for all $i \in I$. Note that $H_i = H_{\{i\}}$, so that all original hypotheses, known as *elementary hypotheses*, are also intersection hypotheses. All are first tested with a valid α-level test which is called the *local test*. This local test should have the property that it rejects H_T, the intersection of all true hypotheses, with probability at most α.

An intersection hypothesis H_I is subsequently rejected by the closed testing procedure if and only if all intersection hypotheses H_J with $I \subseteq J$ are rejected by the local test. By this constraint the closed testing procedure controls the familywise error rate for all intersection hypotheses H_I: the probability that no true H_I is rejected is at least $1 - \alpha$. To see this, note that H_I is true if and only if $I \subseteq T$. Therefore, if H_T is not rejected, no true H_I is rejected by the closed testing procedure. Since H_T is not rejected with probability at least $1 - \alpha$, we have the required result. In particular, therefore, the closed testing procedure controls the FWER for all elementary hypotheses. FWER control for the other intersection hypotheses, however, is also guaranteed.

Based on the result of a closed testing procedure, Goeman and Solari (2011) showed that the required bound with property (1) can be calculated as

$$t_\alpha(S) = \max\{|I| \colon I \subseteq S, H_I \text{ not rejected}\}, \tag{2}$$

where 'rejected' means rejection by the closed testing procedure. This result holds for any closed testing procedure. In the rest of the paper we will study computation of $t_\alpha(S)$ and the properties of this bound for the special case that the local test is a Simes test.

Goeman and Solari (2011) propose not only to report the FDP bounds $q_\alpha(S)$ for usual values $\alpha = 0.05$ or $\alpha = 0.10$, but also to use $\alpha = 0.5$ to get a point estimate of

the FDP. By (1) this point estimate is median unbiased simultaneously over all S. We denote $\hat{q}(S) = q_{1/2}(S)$, and analogous for t and d.

3 Simes and Hommel

Closed testing with Simes local tests is the basis of Hommel's FWER-controlling procedure. Simes test is based on the p-values of elementary hypotheses. We suppose that we have raw (unadjusted) p-values $p_1, \ldots, p_m$, one for each hypothesis $H_1, \ldots H_m$. Without loss of generality we will assume that $p_1 \leq \ldots \leq p_m$.

For closed testing with Simes local tests to be valid, we must assume that the Simes inequality holds for the p-values corresponding to true hypotheses. This means that

$$P\left(\bigcup_{i \in \{1,\ldots,|T|\}} \left\{ p_{(i:T)} \leq \frac{i\alpha}{|T|} \right\} \right) \leq \alpha, \tag{3}$$

where for any $I \subseteq \{1, \ldots, m\}$ and $i \in \{1, \ldots, |I|\}$, we define $p_{(i:I)}$ as the ith smallest p-value among the multiset $\{p_i \colon i \in I\}$. There is extensive and growing literature on the conditions under which this assumption holds, which we will not revisit here (Benjamini and Yekutieli, 2001; Rødland, 2006; Sarkar, 2008; Finner et al., 2014). The assumption (3) is also necessary for the validity of the procedure of Benjamini and Hochberg (1995).

Based on the Simes inequality we can construct Simes local tests: an intersection hypothesis H_I is rejected by that test if and only if there is at least one $i \in \{1, \ldots, |I|\}$ for which

$$p_{(i:I)} \leq \frac{i\alpha}{|I|}. \tag{4}$$

This has the required property that it rejects H_T with probability at most α by assumption (3). For this paper it will often be convenient to write equations in terms of the ordered p-values of the complete set, so we note that the Simes test is equivalent to rejecting H_I if and only if there is at least one $i \in I$ for which

$$p_i \leq \frac{\mathrm{rk}(i : I)\alpha}{|I|},$$

where $\mathrm{rk}(i : I) = |\{j \in I : j \leq i\}|$ is the rank of p-value p_i within the multiset of p-values $\{p_i \colon i \in I\}$, with ties broken arbitrarily.

Naive application of the closed testing procedure requires 2^m local tests to be performed, which severely limits the usefulness of the procedure in large problems. For this reason *shortcuts* have been developed, which are algorithms that limit the computation time of the closed testing procedure. Shortcuts come in two flavors: exact shortcuts, which lead to the same rejections as the closed testing procedure, and approximate or conservative shortcuts, which lead to at most the same rejections as the closed testing procedure, but possibly fewer. Shortcuts always apply to specific local

tests; there are no general shortcuts for the closed testing procedure. Hommel's method is an exact shortcut to the closed testing procedure with Simes local tests, but restricts itself to reporting only the rejections of elementary hypotheses. In contrast, Hochberg's method is a faster conservative shortcut for the same closed testing procedure, also exclusively focused on the elementary hypotheses.

Hommel (1988) proved that H_i is rejected by the closed testing procedure if and only if

$$h(\alpha)p_i \leq \alpha, \tag{5}$$

where

$$h(\alpha) = \max\left\{i \in \{0, \ldots, m\} : ip_{m-i+j} > j\alpha, \text{ for } j = 1, \ldots, i\right\}. \tag{6}$$

Note that (5) is analogous to Bonferroni, except that m is replaced by $h(\alpha) \leq m$. A graphical illustration of $h(\alpha)$ is given in Figure 1. The quantities z and b will be introduced later.

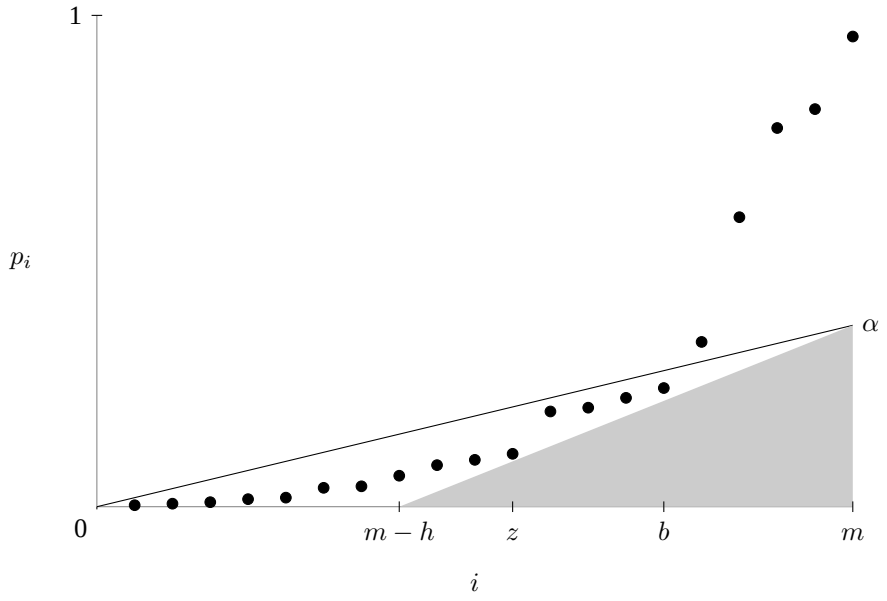

Figure 1: Graphical illustration of $h = h(\alpha)$, z and b. The points are the p-values plotted against their rank. The value of h is the largest such that the grey triangle does not contain any of the points. The value of z is the smallest index i such that $(i+1, p_i)$ is inside the triangle. The value of b is the index of the largest point below the line. The relationship $m - h \leq z \leq b$ is clear.

Hommel's simple rule allows to quickly determine which hypotheses are rejected at a certain α. It is often also of interest to calculate *adjusted p-values* (Wright, 1992).

The adjusted p-value $\tilde{p}_i$ of hypothesis H_i is defined as the smallest α for which the multiple testing procedure rejects H_i. For a closed testing procedure, the adjusted p-value for an intersection hypothesis H_I is the minimum value of α for which all intersection hypotheses H_J with $I \subseteq J$ are rejected by the local test. Calculation of these adjusted p-values for the elementary hypotheses is easy from Hommel's rule (5) if $h(\alpha)$ is known as a function of α. These adjusted p-values are never larger than the maximal p-value p_m.

Hommel's rule (5) is exclusively targeted toward inference on elementary hypotheses. We will extend Hommel's procedure to other intersection hypotheses below to be able to make the FDP statements we aim for. First, however, we must look more deeply into the function $h(\alpha)$, which will remain crucial throughout this paper.

4 The function $h(\alpha)$

The most important quantity in Hommel's procedure is the function $h(\alpha)$. In this section we will study the properties of this function, and show how it can be calculated for all α in only $O(m \log(m))$ time. Since by (5), adjusted p-values of Hommel's method are a simple function of $h(\alpha)$, we can also calculate these in $O(m \log(m))$ time.

The defining property of $h(\alpha)$ is given in Lemma 1. By this lemma $h(\alpha)$ is the largest size $|I|$ of an intersection hypothesis H_I that cannot be rejected by the closed testing procedure at level α. This result was already known to Hommel. We give the proof for completeness.

Lemma 1. *The closed testing procedure with Simes local tests rejects all hypotheses H_I with $|I| > h(\alpha)$. It fails to reject H_K with $|K| = h(\alpha)$, given by $K = \{m - h(\alpha) + 1, \ldots, m\}$.*

Proof. First, for any I with $|I| > h(\alpha)$, consider the 'worst case'

$$J = \{m - |I| + 1, \ldots, m\}.$$

By definition of $h(\alpha)$ as a maximum there is an $i \in \{1, \ldots, |I|\}$ such that

$$|I| p_{(i:I)} \leq |I| p_{(i:J)} = |I| p_{m-|I|+i} \leq i\alpha,$$

which shows that H_I is rejected by the Simes test. Since all H_J with $J \supset I$ are also rejected by the Simes test according to the same argument, H_I is also rejected by the closed testing procedure. This proves the first statement. Next, by the definition of $h(\alpha)$, we have for all $k = 1, \ldots, h(\alpha) = |K|$,

$$|K| p_{(k:K)} = |K| p_{m-h(\alpha)+k} = h(\alpha) p_{m-h(\alpha)+k} > k\alpha$$

so that H_K is not rejected by the Simes test. Consequently it is not rejected by the closed testing procedure. This proves the second statement. □

By Lemma 1 and equations (2) and (1), $h(\alpha)$ is a confidence bound for the number of true null hypotheses $|T|$. The property that $\mathrm{P}(|T| \geq h(\alpha)) \geq 1 - \alpha$ follows immediately from the FWER-control of the closed testing procedure.

The function $h(\alpha)$ is a right-continuous step function on its domain $[0, 1]$, since it is non-increasing and takes integer values in $\{0, \ldots, m\}$. To describe $h(\alpha)$ it suffices, therefore, to find the jumps of the function. We have $h(1) = 0$, and $h(0)$ is the number of non-zero p-values. Let α_i, for $i = 1, 2, \ldots$, be such that $h(\alpha) \geq i$ for all $\alpha < \alpha_i$ and $h(\alpha_i) < i$. By Lemma 1, α_i for $1 \leq i \leq m$ is the lowest α-level at which the hypothesis H_I with $I = \{m - i + 1, \ldots, m\}$, the 'worst case hypothesis of size i', can still be rejected. By (4) this point is given by

$$\alpha_i = \min_{k \in \{1,\ldots,i\}} \frac{i}{k} \cdot p_{m-i+k} = i \cdot \min_{k \in \{1,\ldots,i\}} \frac{p_{m-i+k}}{k}. \tag{7}$$

For $i > m$ we have $\alpha_i = 0$. Tied values of α_i indicate jumps of $h(\alpha)$ larger than 1. However, it can be shown that we can only have $\alpha_i = \alpha_j$ for $i \neq j$ if $\alpha_i = \alpha_j = 0$ or $\alpha_i = \alpha_j = p_m$. This is intuitive from Figure 2, where we illustrate the function $h(\alpha)$ and its jumps.

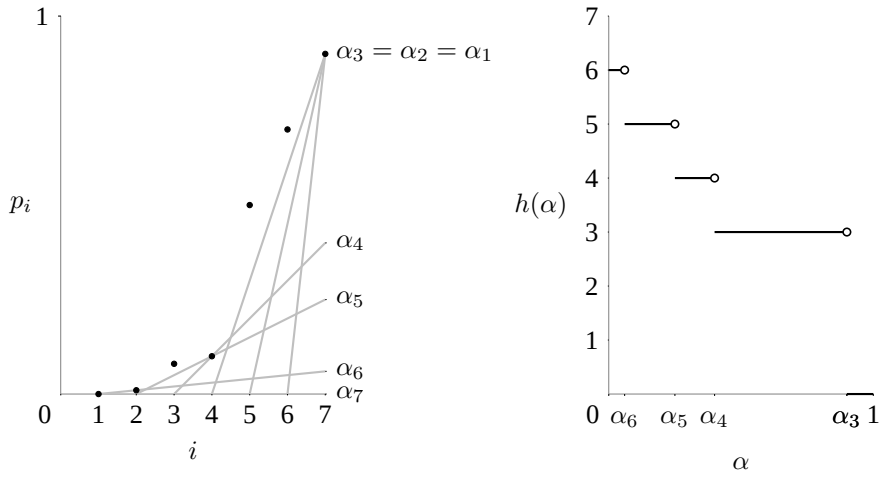

Figure 2: On the left-hand side an example of the construction of α_i, $i = 1, \ldots, m$, for $m = 7$ and $p_1 = 0$, $p_2 = 0.01$, $p_3 = 0.08$, $p_4 = 0.1$, $p_5 = 0.7$, $p_6 = 0.8$, $p_7 = 0.9$. The points are the p-values plotted against their rank. The value of α_i is the largest α such that the line from $(m - i, 0)$ to (m, α) has no points below the line. For the given p-values we find $\alpha_1 = \alpha_2 = \alpha_3 = p_7 = 0.9$, $\alpha_4 = 4p_4 = 0.4$, $\alpha_5 = 5p_4/2 = 0.25$, $\alpha_6 = 6p_2 = 0.06$, $\alpha_7 = 7p_1 = 0$. The right-hand side is the resulting step function $h(\alpha)$.

Finding $h(\alpha)$ using (7) requires computing m minima. These calculations are

equivalent to finding the minimum of each column of the matrix

$$M = \begin{pmatrix} p_1 & & & & & \\ \frac{p_2}{2} & p_2 & & & & \\ \frac{p_3}{3} & \frac{p_3}{2} & p_3 & & & \\ \frac{p_4}{4} & \frac{p_4}{3} & \frac{p_4}{2} & p_4 & & \\ \vdots & \vdots & \vdots & \vdots & \ddots & \\ \frac{p_m}{m} & \frac{p_m}{m-1} & \frac{p_m}{m-2} & \frac{p_m}{m-3} & \cdots & p_m \end{pmatrix}, \tag{8}$$

a calculation that normally takes $\Theta(m^2)$ time. This is how Hommel's procedure is currently implemented, e.g. in the R-function `p.adjust`. However, the complexity of this calculation can be reduced to $O(m \log(m))$ by exploiting the following Lemma.

Lemma 2. For matrix M of the form (8) with $p_1 \leq \ldots \leq p_m$, if a minimum of column k is in row i, a minimum of column $k' < k$ will be in row $j \leq i$; a minimum in column $k' > k$ will be in row $j \geq i$.

Proof. Suppose a minimum in column k is found on row i, which means the minimum is given by $\frac{p_i}{i-k+1}$. Being a minimum means that

$$\forall j \neq i, j \in \{k, \ldots, m\} : \frac{p_i}{i-k+1} \leq \frac{p_j}{j-k+1}. \tag{9}$$

For the first statement of the lemma it suffices to show that for a column $k' < k$, we have

$$\forall j > i \colon \frac{p_i}{i-k'+1} \leq \frac{p_j}{j-k'+1}.$$

To show that this holds, take $j > i$ and multiply both sides of equation (9) by $(i-k+1)(j-k+1)$ which gives

$$p_i(j-k+1) \leq p_j(i-k+1).$$

Since $p_i \leq p_j$ for $i < j$ and $k' < k$, we have $(k-k')p_i \leq (k-k')p_j$. Adding this inequality gives

$$p_i(j-k'+1) \leq p_j(i-k'+1).$$

Dividing both sides by $(j-k'+1)(i-k'+1)$ now gives the desired result.

The second statement of the lemma follows immediately from the first by contradiction. □

Lemma 2 can be used to find the minimum in every column of M in $O(m \log(m))$ steps by dividing the matrix in submatrices in every step and always computing the minimum in the middle column of a submatrix, as follows. We start by finding the minimum of the middle column k, which we will assume to be at row i. This will take

$O(m)$ steps. The search for minima can now be naturally restricted to two submatrices; one consisting only of columns $1, \ldots, k-1$ and rows $1, \ldots, i$, and the second consisting only of columns $k+1, \ldots, m$ and rows $i, \ldots, m$. By Lemma 2 no minima can be found outside these submatrices. Next, in both submatrices we find the minimum of their middle column. This time, calculating both minima simultaneously will take $m+1 = O(m)$ steps. Again, each submatrix can be divided into two new submatrices using Lemma 2, and we iterate. In each step, twice the number of minima compared to the previous step are calculated, and exhausting all m columns will thus take $\lceil \log_2(m) \rceil$ iterations. Since in each step all minima can be simultaneously calculated in $O(m)$ time, the complexity of finding all minima of the columns of M is $O(m \log(m))$. This is the complexity for finding $\alpha_1, \ldots, \alpha_m$, and therefore $h(\alpha)$. The algorithm is summarized in Procedure 1, below.

Procedure 1. *Do* $\mathrm{findjumps}(1, m, 1, m)$*, where* $\mathbf{findjumps}(\underline{r}, \bar{r}, \underline{c}, \bar{c})$*:*

- *Set* $j = \lfloor \frac{1}{2}(\underline{c} + \bar{c}) \rfloor$.
- *Find any* $\max(\underline{r}, j) \leq i \leq \bar{r}$*, that minimizes* $M_{ij} = \frac{p_i}{i-j+1}$.
- *Set* $\alpha_{m-j+1} = (m-j+1)M_{ij}$.
- **if** $j > \underline{c}$ **do** $\mathrm{findjumps}(\underline{r}, i, \underline{c}, j-1)$.
- **if** $j < \bar{c}$ **do** $\mathrm{findjumps}(i, \bar{r}, j+1, \bar{c})$.

end.

From $h(\alpha)$ we can subsequently calculate the adjusted p-value $\tilde{p}_i$ for hypothesis H_i in Hommel's procedure by finding the minimum α for which $h(\alpha)p_i \leq \alpha$. This adjusted p-value is given by

$$\tilde{p}_i = \min(t_i p_i, \alpha_{t_i}),$$

where

$$t_i = \max\{1 \leq j \leq m+1 \colon (j-1)p_i \leq \alpha_j\}. \tag{10}$$

as proven below in Lemma 4. We can use Procedure 2 to find $t_1, \ldots, t_m$. To see that the procedure is valid, note that for each i the procedure finds the largest value of $1 \leq j \leq t_{i-1}$ such that $(j-1)p_i \leq \alpha_j$. Since $p_i \geq p_{i-1}$ we have $t_i \leq t_{i-1}$, so this is equivalent to (10).

Procedure 2. *Initialize* $i = 1$ *and* $j = m+1$*. Repeat*

- *If* $(j-1)p_i \leq \alpha_j$ *set* $t_i = j$ *and increment* i *by 1;*
- *Otherwise, decrease* j *by 1;*
- *If* $i > m$*, stop.*

It is clear from this result that the adjusted p-values of all m elementary hypotheses can be calculated simultaneously in $O(m)$ time after $h(\alpha)$ has been found. Taken together with the calculation of $h(\alpha)$, we can find all m of Hommel's adjusted p-values in $O(m \log(m))$ time. This is a large gain relative to the current $\Theta(m^2)$ time algorithm as implemented e.g. in the `R`-function `p.adjust`. Since sorting of the p-values also takes $\Theta(m \log(m))$ time, the computational complexity of finding sorted adjusted p-values for Hommel's and Hochberg's method thus becomes identical. This removes much of the attractiveness of Hochberg's method relative to Hommel's.

For the rest of this paper we will suppress the dependence on α of $h = h(\alpha)$ and most other quantities in the notation.

5 Non-consonant rejections

The literature on closed testing makes a distinction between *consonant* and *non-consonant* procedures (Henning and Westfall, 2015; Brannath and Bretz, 2010; Bittman et al., 2009; Gou et al., 2014). A consonant closed testing procedure has the property that for every I rejection of H_I by the closed testing procedure implies rejection of H_i for at least one $i \in I$. Procedures without this property are non-consonant. From the perspective of FWER control, consonance is a desirable property. From the perspective of FDP control, this is not necessarily the case (Goeman and Solari, 2011). We define a *non-consonant rejection* as a rejection of H_I for which all H_i, $i \in I$ are not rejected.

Hommel's procedure is not consonant. A simple example serves to demonstrate this. Take for example $m = 4$, with $\alpha/2 < p_1, p_2, p_3 \leq 2\alpha/3$, and $p_4 > \alpha$. Then it is easy to verify that the Simes test rejects H_I for all I except $I = \{1, 4\}$, $\{2, 4\}$, $\{3, 4\}$ and $\{4\}$. Consequently, the closed testing procedure rejects H_I with $I = \{1, 2\}$, $\{1, 3\}$ and $\{2, 3\}$, and all H_I with $|I| > 2$, but fails to reject H_1, H_2, H_3 as well as those intersection hypotheses for which the Simes test failed. The end result, illustrated in Figure 3, is clearly non-consonant. This may seem like a contrived example. Indeed, non-consonant rejections in Hommel's procedure are rare if m is small. They occur very frequently if m is large, however, as we shall see in Section 9.

Motivated by a flow chart of Hommel (1986, Fig. 3), we can directly determine whether a hypothesis H_I can be rejected within the closed testing procedure, generalizing Hommel's rule (5) to H_I with $|I| > 1$. We formulate this generalization as Lemma 3.

Lemma 3. *The closed testing procedure with Simes local tests rejects H_I if and only if there exists some $i \in \{1, \ldots, |I|\}$ such that*

$$hp_{(i:I)} \leq i\alpha.$$

This condition is very similar to the condition (4) for rejection of H_I by the Simes test, but with $|I|$ replaced by h. As in that case, we may use the equivalent condition

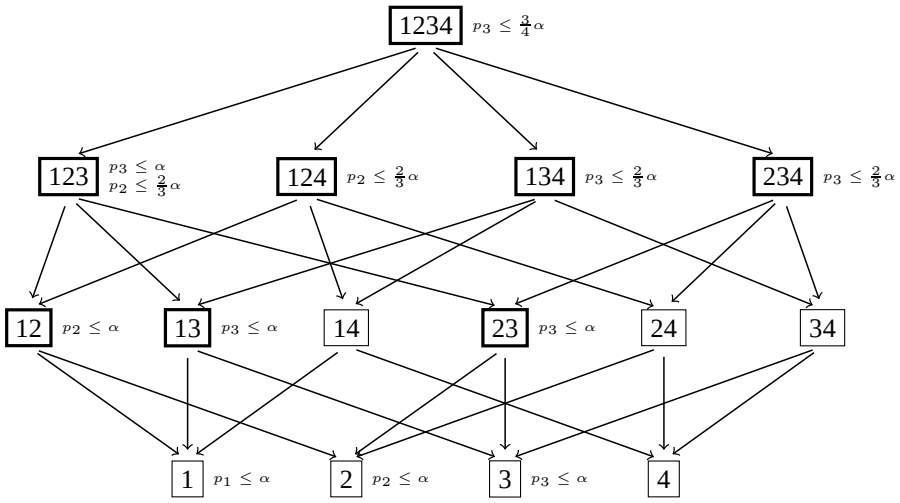

Figure 3: Example of non-consonant rejections in Hommel's procedure based on $\alpha/2 < p_1 \le p_2 \le p_3 \le 2\alpha/3$, and $p_4 > \alpha$. Nodes in the graph represent intersection hypotheses, labeled by their corresponding index set. Edges represent subset relationships. Hypotheses rejected by the Simes test have the reason(s) for their rejection given next to them. Hypotheses rejected by the closed testing procedure are marked in bold. A non-consonant rejection is e.g. $H_{\{1,3\}}$, since it is rejected but H_1 and H_3 are not.

that there exists some $i \in I$ such that

$$hp_i \le \mathrm{rk}(i:I)\alpha.$$

We can also reformulate the statement of Lemma 3 as saying that H_I is rejected if and only if there is some $u \ge 1$ for which

$$\sum_{i \in I} \mathbb{1}\{hp_i \le u\alpha\} \ge u, \tag{11}$$

where $\mathbb{1}\{\cdot\}$ is an indicator function. This reformulation will turn out to be helpful once we study confidence bounds, which we do in the next section.

Proof. Consider first the case $|I| > h$. By Lemma 1, H_I for all such I is rejected by the closed testing procedure, and consequently by the Simes test. For all such H_I, there is an $i \in I$ such that

$$hp_i \le |I| p_i \le \mathrm{rk}(i:I)\alpha$$

Now look at the case $|I| \le h$. First, suppose that there exists some $i \in I$ with $hp_i \le \mathrm{rk}(i:I)\alpha$. We show for all $J \supseteq I$ that H_J is rejected by the Simes test. By

Lemma 1, all H_J with $|J| > h$ are rejected by the Simes test, so we need to consider only $|J| \leq h$. For such J, we have

$$|J|p_i \leq hp_i \leq \mathrm{rk}(i:I)\alpha \leq \mathrm{rk}(i:J)\alpha,$$

which proves that H_J is rejected by the Simes test. Since all H_J with $J \supseteq I$ are rejected by the Simes test, H_I is rejected by the closed testing procedure.

Next, suppose that there is no $i \in I$ with $hp_i \leq \mathrm{rk}(i:I)\alpha$. Consider the set $J \supseteq I$ constructed by augmenting I with $m, m-1, \ldots$, until $|J| = h$. Then $J = I \cup L$ for $L = \{m-k+1, \ldots, m\}$ with $0 \leq k \leq h$. If $j \in I \setminus L$, by the assumption we have

$$|J|p_j \geq hp_j \geq \mathrm{rk}(j:I)\alpha = \mathrm{rk}(j:J)\alpha.$$

If $j \in L$, since $L \subseteq K = \{m-h+1, \ldots, m\}$, we have by Lemma 1,

$$|J|p_j = |K|p_j > \mathrm{rk}(j:K)\alpha = \mathrm{rk}(j:J)\alpha.$$

Consequently, $|J|p_j > \mathrm{rk}(j:J)\alpha$ for all $j \in J$. Therefore H_J is not rejected by the Simes test. Since $J \supseteq I$, H_I is not rejected by the closed testing procedure. □

We can use Lemma 3 to calculate an adjusted p-value for H_I, defined as above as the smallest $\alpha \geq 0$ that allows rejection of H_I by the closed testing procedure.

Lemma 4. *The adjusted p-value of H_I is given by*

$$\tilde{p}_I = \min(tp_I, \alpha_t),$$

where

$$t = \max\{1 \leq j \leq m+1 \colon (j-1)p_I \leq \alpha_j\}.$$

and $p_I = \min_{1 \leq i \leq |I|} p_{(i:I)}/i$.

Proof. We first remark that, by Lemma 3, H_I is rejected by the closed testing procedure if and only if $hp_I \leq \alpha$.

Let $\alpha \geq \tilde{p}_I$. Then $\alpha \geq \alpha_t$ or $\alpha \geq tp_I$. If $\alpha \geq \alpha_t$, then $h < t$. By definition of t we have $hp_I \leq (t-1)p_I \leq \alpha_t \leq \alpha$, which implies that H_I is rejected by the closed testing procedure. If $\alpha \geq tp_I$, then by definition of t if $t \leq m$ and trivially if $t = m+1$, we have $\alpha \geq tp_I \geq \alpha_{t+1}$. Therefore $h \leq t$, so that $hp_I \leq tp_I \leq \alpha$, which implies that H_I is rejected by the closed testing procedure.

Let $\alpha < \tilde{p}_I$. Then $\alpha < tp_I$ and $\alpha < \alpha_t$. The latter implies that $h \geq t$, so that $hp_I \geq tp_I > \alpha$. This implies that H_I is not rejected by the closed testing procedure. □

Using Lemma 4 a single adjusted p-value $\tilde{p}_I$ can be calculated from p_I using bisection in $O(\log(m))$ time, since α_j is non-increasing in j. Alternatively, if sorted values of p_I for $O(n)$ intersection hypotheses are available, Procedure 2 can be used to calculate adjusted p-values in $O(n+m)$ time.

6 The confidence bound

We now turn to calculation of $t_\alpha(S)$ as given in (2), the $(1-\alpha)$-lower confidence bound to the number of true hypotheses in S. Let $\mathbb{N} = \{1, 2, \ldots\}$ as usual. We prove the following lemma.

Lemma 5.

1. *We have*
$$t_\alpha(S) = |S| - \max_{u \in \mathbb{N}} \left\{ 1 - u + \sum_{i \in S} \mathbb{1}\{hp_i \le u\alpha\} \right\}.$$

2. *The maximum is attained for some* $1 \le u \le z - m + h + 1$*, where*
$$z = \begin{cases} 0, & \text{if } h = m; \\ \min\left\{ m - h \le i \le m \colon hp_i \le (i - m + h + 1)\alpha \right\}, & \text{otherwise.} \end{cases} \tag{12}$$
If S *is nonempty, we may even take*
$$1 \le u \le \min\left\{ z - m + h + 1, |S|, \min\{v \in \mathbb{N} \colon h \max_{i \in S} p_i \le v\alpha\} \right\}.$$

A graphical illustration of z is given in Figure 1.

Proof. Let $d = \max_{u \in \mathbb{N}} g(u)$, where
$$g(u) = 1 - u + \sum_{i \in S} \mathbb{1}\{hp_i \le u\alpha\}.$$
and let $\tilde{u}$ be the smallest location of the maximum, so that
$$1 - k + \sum_{i \in S} \mathbb{1}\{hp_i \le \tilde{u}\alpha\} = d.$$
Note that $0 \le d \le |S|$ and $1 \le \tilde{u} \le |S| - d + 1$.

First, suppose $d > 0$ and take any $I \subseteq S$ such that $|I| \ge |S| - d + 1$. We have
$$\begin{aligned} \sum_{i \in I} \mathbb{1}\{hp_i \le \tilde{u}\alpha\} &= |I| - \sum_{i \in I} \mathbb{1}\{hp_i > \tilde{u}\alpha\} \\ &\ge |I| - \sum_{i \in S} \mathbb{1}\{hp_i > \tilde{u}\alpha\} \\ &= |I| - |S| + \sum_{i \in S} \mathbb{1}\{hp_i \le \tilde{u}\alpha\} \\ &\ge (1 - d) + (d + \tilde{u} - 1) \;=\; \tilde{u}, \end{aligned}$$
which shows by Lemma 3, reformulated in (11), that H_I is rejected by the closed testing procedure. Therefore, $t_\alpha(S) \le |S| - d$. The same trivially holds if $d = 0$.

Next, consider the set $K \subseteq S$ with $|K| = |S| - d$ constructed by discarding the smallest d elements from S. We show that H_K is not rejected by the closed testing procedure. We proceed by contradiction. Suppose H_K is rejected. Then by Lemma 3 there is some $u \in \mathbb{N}$ such that $\sum_{i \in K} \mathbb{1}\{hp_i \leq u\alpha\} \geq u$. Then $\sum_{i \in S} \mathbb{1}\{hp_i \leq u\alpha\} \geq d + u$ by construction of K. So

$$1 - u + \sum_{i \in S} \mathbb{1}\{hp_i \leq u\alpha\} > d,$$

but this contradicts the definition of d as a maximum. Therefore H_K is not rejected, so that $t_\alpha(S) \geq |S| - d$.

Together, this proves that $t_\alpha(S) = |S| - d$. This proves Part 1.

To prove Part 2, let us first verify that z is well-defined. Since h is a maximum, we have for $h < m$ that $hp_{m-h+i-1} \leq (h+1)p_{m-h+i-1} \leq i\alpha$ for some $1 \leq i \leq h+1$, so we have $z = m-h+i-1$ for the smallest such i. Pick any $z-m+h+1 \leq u \leq m$.

Next, we will show that $g(z - m + h + 1) \geq g(u)$, so $\tilde{u} \leq z - m + h + 1$. By definition of h, we have $hp_{m-h+i} > i\alpha$ for all $1 \leq i \leq h$, which becomes, after changing variables, $hp_i > (i - m + h)\alpha$ for all $m - h + 1 \leq i \leq m$. In particular, if $h = m$, then $g(u) \leq 0 \leq g(1)$ because $\sum_{i \in S} \mathbb{1}\{hp_i \leq u\alpha\} \leq u - 1$, so assume $h < m$. Consider any $1 \leq i \leq m$. If $i \leq z$, then $hp_i \leq hp_z \leq (z - m + h + 1)\alpha$ by definition of z. If $i \geq u+m-h > z$, then $i \geq m-h+1$, so $hp_i > (i-m+h)\alpha \geq u\alpha$ by definition of z. Hence,

$$\sum_{i \in S} \mathbb{1}\{(z - m + h + 1)\alpha < hp_i \leq u\alpha\} \leq u + m - h - z - 1,$$

from which the conclusion follows.

If S is non-empty, obviously $g(u)$ is decreasing after

$$u = \min\{v \in \mathbb{N} \colon h \max_{i \in S} p_i \leq v\alpha\},$$

so $\tilde{u} \leq u$. For $u > |S| \geq 1$, we have $g(u) \leq 0 \leq g(1)$, so $\tilde{u} \leq |S|$. □

Clearly, using Lemma 5 the bound $t_\alpha(S)$ for any S can be calculated in linear time in $|S|$ once h has been calculated. We can use z to stop the calculation earlier, but computation time remains linear in $|S|$. Calculation of z takes linear time in m, but since it does not depend on S we need to do it only once, like with h.

In the rest of this section we will go one step further, and calculate $t_\alpha(S_1), \ldots, t_\alpha(S_k)$ for an incremental sequence of sets $S_1 \subset \ldots \subset S_k$ with $|S_i| = i$. We show that the confidence bound for all these sets together can be calculated together in linear time in k (again after h has been calculated).

A sequence $S_1 \subset \ldots \subset S_k$ with $|S_i| = i$ may arise when a user is interested in the top $1, 2, \ldots, k$ hypotheses according to some criterion or combination of criteria. For example, a user may want to look at the top $1, 2, \ldots, k$ estimated effect sizes among

the k hypotheses with p-values smaller than some cut-off. Using the algorithm below, the user can calculate the number or proportion of false positive findings for any such combination of ordering and selection criterion. Remember that the ordering and the selection criterion may be chosen post hoc.

For this algorithm we reformulate the problem in terms of categories. By Lemma 3, to decide whether an intersection hypothesis can be rejected at a fixed α we do not need the exact values of p_i, $i \in I$, but we only need to know how many integer multiples of α/h (if $h > 0$) they make. In general, we define the *category* of a p-value p as

$$\mathrm{c}(p) = \min\big\{u \in \mathbb{N}\colon hp \le u\alpha\big\}. \tag{13}$$

Let c_i, $i = 1, \ldots, k$ be the category corresponding to the p-value with index in S_i but not S_{i-1}. Then we have a (not necessarily increasing) sequence $\mathbf{c} = (c_1, c_2, \ldots, c_k)$, $k \le m$, of categories, and by Lemma 5 we have

$$t_\alpha(S_i) = i - \max_{u \in \{1,\ldots,v\}} \left\{1 - u + \sum_{j=1}^{i} \mathbb{1}\{c_j \le u\}\right\}. \tag{14}$$

where $v = \min(k, z - m + h + 1, \max_j c_j)$. Although we could sometimes take v smaller if we let it depend on i, naive calculation of $t_\alpha(S_i)$ for all $i = 1, \ldots, k$ will generally take $\Theta(kv)$ time unless the categories are ordered.

However, a procedure that can be implemented in linear time is given below. It uses partitioning of the index set $\{1, \ldots, v\}$ into disjoint intervals. The intervals in the partitioning are naturally ordered by the lowest index they contain. For a partitioning π, denote its jth smallest interval by π_j.

Procedure 3. *Initialize $d_0 = 0$ and $\pi = \big(\{1\}, \{2\}, \ldots, \{v\}\big)$. For $i = 1, \ldots, k$,*

- *if $c_i \in \pi_1$, set $d_i = d_{i-1} + 1$ and leave π unchanged.*
- *if $c_i \in \pi_q$ for $q > 1$, set $d_i = d_{i-1}$ and merge π_q with π_{q-1}.*
- *if $c_i > v$, set $d_i = d_{i-1}$ and leave π unchanged.*

Note that because nothing happens if $c_i > v$, categories greater than v may be removed from $\mathbf{c}$ beforehand. We come back to this in Section 7.

An example of Procedure 3 for $\mathbf{c} = (4, 2, 2, 3, 7, 3)$ with $v = 5$ is given in Table 1. We prove the validity of this procedure in Lemma 6.

Lemma 6. *For $i = 1, \ldots, k$, let d_i be defined by Procedure 3 and $t_\alpha(S_i)$ by (14), then*

$$t_\alpha(S_i) = i - d_i.$$

Proof. Let $g(u) = 1 - u + \sum_{j \le i} \mathbb{1}\{c_j \le u\}$. Suppose that after i iterations the partition π has $\pi_q = \{a_q, \ldots, b_q\}$. The number of categories in each interval is easy to count:

$$\sum_{j \le i} \mathbb{1}\{c_j \in \pi_q\} = \begin{cases} b_1 - a_1 + d_i, & \text{if } q = 1; \\ b_q - a_q, & \text{otherwise.} \end{cases}$$

Table 1: A sample run of the Procedure 3 to calculate d_i, $i = 1, \ldots, k$, on a sequence of categories $\mathbf{c} = (4, 2, 2, 3, 7, 3)$ with $v = 5$. In each step i, either the partition is updated by merging the interval that contains c_i with its precursor, or d_{i-1} is incremented by 1 if no such precursor exists. Changes relative to the previous step are given in bold.

i	c_i	d_i	π_1	π_2	π_3	π_4	π_5	comment
–	–	0	$\{1\}$	$\{2\}$	$\{3\}$	$\{4\}$	$\{5\}$	initialization
1	4	0	$\{1\}$	$\{2\}$	$\mathbf{\{3,4\}}$	$\{5\}$		since $c_1 \in \pi_4$
2	2	0	$\mathbf{\{1,2\}}$	$\{3,4\}$	$\{5\}$			since $c_2 \in \pi_2$
3	2	**1**	$\{1,2\}$	$\{3,4\}$	$\{5\}$			since $c_3 \in \pi_1$
4	3	1	$\mathbf{\{1,2,3,4\}}$	$\{5\}$				since $c_4 \in \pi_2$
5	7	1	$\{1,2,3,4\}$	$\{5\}$				since $c_5 > v$
6	3	**2**	$\{1,2,3,4\}$	$\{5\}$				since $c_6 \in \pi_1$

We see that $g(b_1) = d_i$.

Pick any $1 \leq u \leq v$. For some p we find $u \in \pi_p$. We count

$$\sum_{j \leq i} \mathbb{1}\{u < c_j \leq b_p\} \geq b_p - u.$$

Therefore,

$$\begin{aligned} \sum_{j \leq i} \mathbb{1}\{c_j \leq u\} &= \sum_{q=1}^{p} \sum_{j \leq i} \mathbb{1}\{c_j \in \pi_q\} - \sum_{j \leq i} \mathbb{1}\{u < c_j \leq b_p\} \\ &\leq d_i + \sum_{q=1}^{p} (b_q - a_q) - (b_p - u) = d_i + p - u. \end{aligned}$$

Hence, $g(u) \leq d_i - p + 1 \leq d_i$. Therefore $\max_{u \in \{1,\ldots,v\}} g(u) = d_i$. The result now follows from Lemma 5. □

The algorithm of Procedure 3 lends itself naturally for implementation with a disjoint-set data structure. Such a data structure allows for efficient queries to which set an element belongs to as well as for taking efficient set unions. Generally, the performance of a disjoint-set data structure for a linear number of operations is not linear (Fredman and Saks, 1989) but nearly linear in time (Tarjan, 1975). In our specific case it is in fact possible to achieve linear time since we only need to merge adjacent intervals instead of arbitrary sets (Gabow and Tarjan, 1985). However, in practice the gain of the specific linear variant over the general nearly linear disjoint-set data structure is not noticeable so we opted for the latter, simpler implementation in the `cherry` package.

7 Concentration of results

Interesting results, of course, mostly happen for the smallest p-values. Two index sets are of particular interest because they play a special role in the results of the closed testing procedure with Simes local tests.

First and foremost is Hommel's rejected set

$$R = \{i\colon hp_i \leq \alpha\}.$$

For this set we have the property that $t_\alpha(S) = 0$ if and only if $S \subseteq R$. This is direct from the relationship between Hommel's procedure and closed testing. By definition of $t_\alpha(S)$ we have $t_\alpha(S) = 0$ if and only if H_i is rejected by the closed testing procedure for all $i \in S$, which happens if and only if $S \subseteq R$. Practically this means that a user who is interested in a subset of the set returned by the FWER-corrected procedure at level α obtains $(1-\alpha)$-confidence of no false positives. This is exactly the same that FWER control guarantees. A user may therefore choose post hoc to control FWER, rather than some more lenient FDP level. A user may opt for this strategy if FWER control already results in an acceptable number of rejections.

Next, consider the *concentration set* Z given by

$$Z = \{1 \ldots, z\},$$

where z is defined in (12). Obviously, $Z \supseteq R$. We call Z the concentration set because of its properties defined in Lemma 7. Remember that $d_\alpha(S) = |S| - t_\alpha(S)$, the $(1-\alpha)$-lower confidence bound on the number of correct selections.

Lemma 7. *We have*

1. $d_\alpha(S) = d_\alpha(S \cap Z)$ *for all* S*;*
2. $d_\alpha(Z) = m - h$*;*
3. Z *is minimal, i.e. if* $d_\alpha(S) = d_\alpha(S \cap Y)$ *for all* S*, then* $Z \subseteq Y$*.*

Proof. The entire lemma is trivial if $h = m$. Assume $h < m$. In that case we remark that by definition of h we have $hp_i > (i - m + h)\alpha$ for all $i \in S \setminus Z$, since $i > z \geq m - h$, so $hp_i > (i - m + h)\alpha \geq (z - m + h + 1)\alpha$. Now $d_\alpha(S) = d_\alpha(S \cap Z)$ follows directly from Lemma 5 since p-values for which $hp_i > (z - m + h + 1)\alpha$ induce no change in $d_\alpha(S)$. This proves Statement 1. We use that statement to obtain $d_\alpha(Z) = d_\alpha(Z \cap \{1, \ldots, m\}) = d_\alpha(\{1, \ldots, m\}) = m - h$, where the final equality holds by definition of d_α and Lemma 1. This proves Statement 2.

Now take any Y that has the property of Statement 3. Consider $Z' = Y \cap Z$. We have

$$m - h = d_\alpha(Z) = d_\alpha(Z') = \max_{u \in \mathbb{N}} \left\{ 1 - u + \sum_{i \in Z'} \mathbb{1}\{hp_i \leq u\alpha\} \right\}. \tag{15}$$

By definition of z as a minimum we have for all $m-h \leq j < z$ that $hp_j > (j-m+h+1)\alpha$, and the same trivially holds for $j < m-h$. Therefore, $\sum_{i \in Z'} \mathbb{1}\{hp_i \leq (u-m+h+1)\alpha\} < u$, for all $u < z$. Changing variables, we obtain $\sum_{i \in Z'} \mathbb{1}\{hp_i \leq u\alpha\} < u+m-h-1$ for all $u < z-m+h+1$, so that $1-u+\sum_{i \in Z'} \mathbb{1}\{hp_i \leq u\alpha\} < m-h$ for all $u < z-m+h+1$. Therefore, and by Lemma 5 the maximum in (15) is attained at $u = z-m+h+1$, and we obtain from (15) that

$$\sum_{i \in Z'} \mathbb{1}\{hp_i \leq (z-m+h+1)\alpha\} = z,$$

so that $|Z'| = |Z|$ and $Z \subseteq Y$. This proves Statement 3. □

Lemma 7 says that all interesting statements from the closed testing procedure are about hypotheses with indices in Z. We are $(1-\alpha)$-confident that there are at least $m-h$ false null hypotheses, and that all these have indices in Z. There is no confidence for any discoveries relating to other hypotheses. For a user looking where the interesting findings of the experiment are, it makes sense to look only at hypotheses with index in Z. If the user restricts to a smaller set, he or she will lose out on some discoveries. In some applications it will make sense to report $t_\alpha(S)$ for sets $S \subseteq Z$ only.

8 Relationship with Benjamini-Hochberg

Like the closed testing procedure studied in this paper, the FDR-controlling procedure of Benjamini and Hochberg (1995), briefly BH, is closely related to the Simes inequality. The conditions for the validity of BH and this closed testing procedure are virtually identical. Moreover, we see that BH rejects at least one hypothesis if and only if hypothesis $H_{\{1,\ldots,m\}}$ is rejected by the closed testing procedure so that the methods make non-trivial statements for exactly the same data sets. The two methods therefore have the same weak FWER control. It is not surprising that there are more close connections between these procedures. We explore a few of these in this section.

The rejection set of BH at level x is given by $B(x) = \{1, \ldots, b(x)\}$, with

$$b(x) = \max\{1 \leq i \leq m \colon mp_i \leq ix\}, \tag{16}$$

taking $b(x) = 0$ and $B(x) = \emptyset$ if the maximum does not exist. We abbreviate $b = b(\alpha)$ and $B = B(\alpha)$.

Interestingly, there is a relationship between b and the quantity z defined in (12) and discussed in the Section 7. This relationship was suggested by the construction in Figure 1. We formulate this as Lemma 8.

Lemma 8. *We have $z \leq b$.*

Proof. Note that $b = z = 0$ if $h = m$. If $h < m$, by definition of h there is a $1 \leq j \leq h+1$ such that $(h+1)p_{j+m-h-1} \leq j\alpha$. Let $z' = j + m - h - 1$. Then

$$hp_{z'} \leq (h+1)p_{z'} \leq (z' - m + h + 1)\alpha,$$

so $z \leq z'$ by definition of z. Now,

$$mp_{z'} = \frac{m(h+1)p_{z'}}{(h+1)} \leq \frac{m(z' - m + h + 1)}{(h+1)z'} z'\alpha = \frac{mz' - m(m-h-1)}{mz' - z'(m-h-1)} z'\alpha \leq z'\alpha,$$

so $z \leq z' \leq b$ by definition of b. □

Combining Lemmas 7 and 8, we have $Z \subseteq B$, so that

$$d_\alpha(B) = m - h. \tag{17}$$

Apparently, there is always a non-trivial number of discoveries at $(1-\alpha)$-confidence in the BH-rejected set at level α. Moreover, all interesting findings of the closed testing procedure at $(1-\alpha)$-confidence are related to hypotheses with indices in a subset of B.

Interestingly, there is also a useful relationship between $B(x)$ for general x and the confidence bounds $q_\alpha(S)$ for the FDP. This relationship follows from Lemma 9.

Lemma 9. *If $q = mp_{(i:S)}/i$, where $i = |S| > 0$, then*

$$m\alpha q_\alpha(S) \leq hq,$$

with equality if and only if $hq = 0$.

Proof. We need to show that all H_J with $J \subseteq S$ and $m\alpha|J| \geq hiq$ are rejected by the closed testing procedure. We have $hp_{(j:J)} \leq hp_{(i:S)} = hiq/m \leq j\alpha$, where $j = |J|$, from which rejection of H_J follows by Lemma 3. □

Bounding $q_\alpha(S)$ by $hq/m\alpha$ for $\alpha > 0$, Lemma 9 improves upon the bound q/α obtained by Lehmann and Romano (2005) using Markov's inequality, which does not have the factor h/m that is stricly smaller than 1 if $q \leq \alpha$. Moreover, the bound of Lehmann and Romano (2005) is not simultaneous over S. The lemma gives a simple $(1-\alpha)$-upper bound for the FDP of a set S, which is especially useful for sets of the form $B(x)$. For such a set, q is exactly the BH-adjusted p-value, so that Lemma 9 gives an immediate upper bound for the FDP confidence bound of such a set. We already know from (17) that if $h < m$ (equivalently $b > 0$), we have $q_\alpha(B) = (b - m + h)/b < 1$, which always gives a non-trivial $(1-\alpha)$-confidence bound for the FDP realized by BH. Additionally from the lemma we see that for all $0 < x \leq 1$ we have $q_\alpha(B(x\alpha)) < x$. By reducing the α-level of BH we may easily guarantee decent tail probabilities of the FDP. Note that through h/m the bound of Lemma 9 automatically incorporates an estimate of $|T|/m$ as adaptive FDR-controlling methods do.

Lemma 9 can also be used to study the behavior of the estimate $\hat{q}(S)$ of the FDP. By the properties of BH, the expected FDP of B is at most α. By Lemma 9 we see that the estimated FDP $\hat{q}(B) = q_{0.5}(B) < 2\alpha$. This is only an upper bound, and the actual value of $\hat{q}(B)$ can be much smaller, even 0 if $B = R$. The upper bound is up to a factor 2 larger than we would hope to see, but even a factor 2 is a surprisingly small price to pay if we realize that $\hat{q}(S)$ is guaranteed to be simultaneously median unbiased not only for $S = B$, but also for $2^m - 1$ other sets S.

9 Scalability of power

A major difference between FWER- and FDR-controlling methods lies in the way the methods scale when m becomes large. The number of rejections of FWER-based methods typically diminishes relative to m, while that of FDR-based methods remains stable. Although constructed from a FWER-controlling closed testing procedure, the FDP confidence statements are, as we have seen in the last section, closely related to the FDR-controlling method of Benjamini and Hochberg. In this section we show that the scaling behavior is also like FDR-controlling methods. Interesting statements from this method do not vanish relative to m as $m \to \infty$.

As a context for studying the power of the method as $m \to \infty$, we use the model of Efron (2012) in which test statistics are sampled, possibly with dependence, from a mixture distribution

$$F(x) = \gamma F_0(x) + (1 - \gamma)F_1(x), \tag{18}$$

where γ is the proportion of true hypotheses, F_0 is the null distribution, and F_1 is the distribution of the test statistic for false null hypotheses. Efron's model allows m to grow while ensuring that later p-values are 'similar' to previous ones. If we reject for high values of the test statistic, P-values based on Efron's model have distribution function $P(x) = 1 - F(F_0^{-1}(1 - x))$. We assume that F_1 and F_0 are regular enough that P is continuous and strictly increasing with $P(0) = 0$ and $P(1) = 1$. Under very mild assumptions on the dependence structure of the p-values we have that the (xm)th ordered p-value for $x \in [0, 1]$ has

$$p_{xm} \xrightarrow{P} P^{-1}(x),$$

where $\xrightarrow{P}$ denotes convergence in probability as $m \to \infty$.

Since h is such a crucial quantity, we start by studying the behavior of h as $m \to \infty$. In this section we assume $0 < \alpha < 1$.

Lemma 10. *For $m \to \infty$, we have*

$$\frac{h}{m} \xrightarrow{P} \min_{0 \le x < P(\alpha)} \frac{x\alpha - \alpha}{P^{-1}(x) - \alpha}.$$

Proof. Let $E = \{p_1 < \alpha \text{ and } p_m > \alpha\}$. Note that $\lim_{m\to\infty} \mathrm{P}(E) = 1$, so since h/m is bounded we may assume E. Similarly, we may assume that all p-values are non-zero.

Starting from (6) and substituting $u = i/m$ and $x = (m + k - i)/m$, we have that h/m is the maximal value of $u \in \{0, 1/m, 2/m, \ldots, 1\}$ such that

$$up_{xm} > (x + u - 1)\alpha \tag{19}$$

for all $x \in \{1-u+1/m, 1-u+2/m, \ldots, 1\}$. Since $up_{xm} > 0$ and $(x+u-1)\alpha \leq 0$ for $x \leq 1 - u$, we may expand the range of x to $x \in \{1/m, 2/m, \ldots, 1\}$. Since $(x + u - 1)\alpha \leq u\alpha$ for all relevant x, u, we can restrict to $p_{xm} < \alpha$, i.e. $x \leq y$, where $y = \max\{x\colon p_{xm} < \alpha\}$. We have that h/m is the maximal value of u in its range such that (19) holds for all $x \in \{1/m, 2/m, \ldots, y\}$. Since $p_{xm} < \alpha$, (19) is equivalent to

$$u < \frac{\alpha - x\alpha}{\alpha - p_{xm}}.$$

Since the right-hand side of this inequality does not depend on u, we have

$$\frac{h}{m} = \frac{1}{m}\Big\lfloor m \cdot \min_{x\in\{1/m,\ldots,y\}} \frac{\alpha - x\alpha}{\alpha - p_{xm}}\Big\rfloor.$$

Letting $m \to \infty$, we have $y \xrightarrow{P} P(\alpha)$ from below, and $p_{xm} \xrightarrow{P} P^{-1}(x)$. By the continuous mapping theorem, we get the result of the lemma, remarking that $(\alpha - x\alpha)/(\alpha - P^{-1}(x)) \to 1$ if $x \downarrow 0$ and $\to \infty$ if $x \uparrow P(\alpha)$, since $P(\alpha) < 1$. □

The limit in Lemma 10 is strictly greater than 0 since $P(\alpha) < 1$. It is strictly smaller than 1 if and only if

$$P^{-1}(x) < x\alpha$$

for at least one $x > 0$. We call a distribution function P for which this holds *Simes-detectable at* α. Chi (2007) studied Simes-detectability under the name criticality in the context of BH. We have

$$\frac{\mathrm{d}}{\mathrm{d}x}P(x) \;=\; \frac{f(F_0^{-1}(1-x))}{f_0(F_0^{-1}(1-x))} \;=\; \gamma + (1-\gamma)\frac{f_1(F_0^{-1}(1-x))}{f_0(F_0^{-1}(1-x))},$$

where f, f_0, and f_1 are the densities corresponding to F, F_0, and F_1, respectively. A sufficient condition for Simes-detectability at all $\alpha > 0$ is that $\gamma < 1$ and $\lim_{x\to\infty} f_1(x)/f_0(x) = \infty$, since in that case $\lim_{x\downarrow 0} \frac{\mathrm{d}}{\mathrm{d}x}P(x) \to \infty$, so that $\lim_{x\downarrow 0} \frac{\mathrm{d}}{\mathrm{d}x}P^{-1}(x) \to 0$. From this condition it is easy to check that Simes-detectability for all $\alpha > 0$ holds true if $\gamma < 1$ and F_0 is standard normal and F_1 is normal with $\sigma^2 > 1$ or with $\sigma^2 = 1$ and $\mu > 0$. In other models, e.g. t and F-distributed test statistics, Simes-detectability α only holds from a certain sample size if $\gamma < 1$. Sample sizes from which Simes-detectability is guaranteed may be calculated from the results of Chi (2007).

FWER-controlling methods generally do not scale well with m. This is easily seen for Hommel's method. Consider the index set $R = \{i : hp_i \leq \alpha\}$ rejected by Hommel's method. For that set, substituting $q = i/m$, and letting $m \to \infty$ we have

$$|R|/m \xrightarrow{P} \sup\{q : P^{-1}(q) \leq 0\} = 0,$$

since $\alpha/h \xrightarrow{P} 0$. If $\gamma < 1$, as m increases, the number of false hypotheses grows proportional to m in the model (18). The number of rejections of Hommel's FWER-controlling method, however, does not grow proportional to m. Therefore, the proportion of rejected hypotheses relative to false hypotheses shrinks to zero. This kind of behavior is typical of FWER-controlling methods.

In contrast, it is well known that FDR controlling methods, such as BH, do scale well with m. Simes-detectability is key here, as was shown by Chi (2007). If P is Simes-detectable at α, the size of the rejection set of BH does not vanish relative to m like Hommel's rejected set, but if $P^{-1}(x) > x\alpha$ for all $x > 0$, we do have $|B|/m \xrightarrow{P} 0$. The borderline case is complex and we do not consider it here.

Lemma 11 proves that the bounds $q_\alpha(S)$ have a similar scalability property to BH.

Lemma 11. *If P is Simes-detectable at level $q\alpha$ for some $q \in (0,1)$, then there is a set S_m with $|S_m|/m \to x$ for some $x \in (0,1)$ such that $q_\alpha(S_m) \xrightarrow{P} q' < q$.*

Proof. If P is Simes-detectable at $q\alpha$, we have $P^{-1}(x) < xq\alpha$ for at least one $x \in (0,1)$. Choose such an x and define $S_m = \{1, \ldots, \lfloor mx \rfloor\}$. Obviously, $|S_m|/m \to x$. By Lemma 5 we have, for S_m of this form,

$$q_\alpha(S_m) = 1 - \max_{u \in \mathbb{N}} \Big\{ \frac{1-u}{\lfloor xm \rfloor} + \frac{1}{\lfloor xm \rfloor} \max\big\{ i \in S_m : hp_i \leq u\alpha \big\} \Big\}.$$

Reparameterizing $y = u/m$, this converges to

$$q_\alpha(S_m) \xrightarrow{P} 1 - \sup_{y \geq 0} \Big\{ -\frac{y}{x} + \frac{1}{x} \min\big\{ x, P\big(\frac{y\alpha}{\bar{h}}\big) \big\} \Big\},$$

where $h/m \xrightarrow{P} \bar{h}$ defines $\bar{h}$. Setting $v = P(y\alpha/\bar{h})$, so that $y = \bar{h}P^{-1}(v)/\alpha$, we get

$$\begin{aligned} q_\alpha(S_m) &\xrightarrow{P} 1 - \sup_{v \geq 0} \Big\{ -\frac{\bar{h}P^{-1}(v)}{x\alpha} + \frac{1}{x}\min(x, v) \Big\} \\ &= 1 - \sup_{0 \leq v \leq x} \Big\{ \frac{v}{x} - \frac{\bar{h}P^{-1}(v)}{x\alpha} \Big\}. \end{aligned} \tag{20}$$

Now we use the assumption of Simes-detectability. Let q' be the limit in (20). Since $P^{-1}(v) \leq P^{-1}(x) < xq\alpha$ for $v \leq x$, we have

$$q' \leq 1 - \sup_{0 \leq v \leq x} \Big\{ \frac{v}{x} - \frac{\bar{h}xq\alpha}{x\alpha} \Big\} = \bar{h}q < q,$$

since $\bar{h} < 1$ because P is Simes-detectable at $q\alpha < \alpha$. □

Lemma 11 shows that, unlike Hommel's FWER-controlling method, FDP confidence statements scale well with m. If Simes-detectability holds and m is large enough, then for any desired FDP $q \in (0, 1)$ we can always select a set that has $(1 - \alpha)$-confidence of an FDP at most q, and the size of this set grows linearly with m. Moreover, we may choose q after seeing the data.

Lemma 11 also shows that there is a huge qualitative difference between inference based on $t_\alpha(S)$ and inference based only on the FWER-controlling procedure of Hommel. If we let $R = \{i : hp_i \leq \alpha\}$ be the rejected set of Hommel's procedure, then $|S \setminus R|$ follows immediately as a FWER-based $(1 - \alpha)$-upper confidence bound for the number of true hypotheses $\tau(S)$. We have $t_\alpha(S) \leq |S \setminus R|$, because if $J \subseteq S$ with $|J| > |S \setminus R|$, we have $i \in J$ for some $i \in R$; since H_i was rejected by the closed testing procedure H_J must be. In Figure 1 we see an example where $1 = t_\alpha(S) < |S \setminus R| = 3$ for $S = \{1, 2, 3\}$. Such a situation becomes increasingly more common when m increases. In fact, lemma 11 shows that the FDP bounds $t_\alpha(S)$ and $|S \setminus R|$ have divergent behavior if $m \to \infty$. For any S_m with $|S_m|/m \to x \in (0, 1)$, we have $|S_m \setminus R|/|S_m| \xrightarrow{P} 1$, because $|R|/m \xrightarrow{P} 0$. In contrast, by Lemma 11, for any desired non-trivial FDP bound, we can find a non-vanishing set for which that bound is controlled.

The qualitative difference in behavior as $m \to \infty$ between Hommel's FWER-controlling procedure and the FDP confidence bounds $t_\alpha(S)$ is exclusively due to the exploitation of non-consonant rejections by the latter. From this result, we see that non-consonant rejections, which are rare in closed testing with Simes if m is small, become more and more frequent as m increases, and that as $m \to \infty$ the number of consonant rejections even vanishes relatively to the number of non-consonant rejections. The step from FWER to FDP control therefore forces a radical reassessment of the value of consonance in closed testing procedures. Non-consonant procedures, which are known to be suboptimal from a FWER perspective (Bittman et al., 2009; Gou et al., 2014), seem to be crucial for powerful FDP inference in large problems.

Interestingly, the result in Lemma 11 is a finite sample result, requiring only that the number of hypotheses $m \to \infty$. The sample size must be large enough compared to the amount of signal to have Simes-detectability at $q\alpha$. But if that finite sample size is reached, the number of discoveries grows linearly with the number of tests, eventually leading to an infinite number of discoveries with $(1-\alpha)$-confidence of an FDP at most q. This is a very powerful result.

10 Discussion

We have explored the construction of $(1 - \alpha)$-upper confidence bounds $q_\alpha(S)$ for the proportion of true hypotheses in an index set S, which are simultaneous for all such sets S. We focused on situations with a large number of hypotheses m and proved good scalability both in terms of computation time and power. We see four main uses

for these confidence bounds, which we detail below.

Flexible FDP control The most classical use would be simply to select the largest index set $S = \{i\colon p_i \leq x\}$ such that $q_\alpha(S)$ or $\hat{q}(S)$ is smaller than some maximal acceptable FDP. This use is similar to classical FDR control except that a median or a tail probability of the FDP is controlled rather than an expectation. A major advantage over regular FDR control is that the level at which the FDP is to be controlled may be chosen after seeing the data without destroying the control property. Users may therefore adapt to the amount of signal in a data set, choosing a stringent FDP or even FWER control if such control allows an acceptable number of rejections, or rather a more lenient FDP if that is needed to obtain sufficient discoveries. Regardless of this prior look at the data, because of the simultaneity of the confidence bound, the FDP of the chosen set remains in control with probability at least $1 - \alpha$. This flexibility is unique among FDR- or FDP-controlling methods. Still, the power of the procedure is acceptable compared to e.g. Benjamini-Hochberg, with $0 \leq \hat{q}(B) < 2\alpha$ for the BH-rejected set B.

FDP control for alternative top lists The most interesting discoveries are not always those with the smallest p-values. Other statistical measures such as estimated effect size may be important to determine the value of a discovery once the p-value is small enough. In transcriptomics, for example, it is common practice to consider only those discoveries that have both a small enough p-value and a large enough effect size estimate. FDR control, however, typically focuses on sets of the form $S = \{i\colon p_i \leq x\}$ only. The FDP confidence bounds in this paper allow FDP control for sets of any construction, and with choice of criteria decided upon after seeing the data. Procedure 3 may be helpful in this situation for monitoring the FDP as a function of the number of top discoveries after they have been ordered according to some criterion of interest (e.g. effect size). It is recommended to restrict this to hypotheses with indices in the concentration set Z.

FDP for non-exchangeable hypotheses Methods that control FDR typically require some assumption of exchangeability among hypotheses. In fact, this exchangeability is implied in the definition of FDR, which allows false inference for one hypothesis to be compensated by correct inference in another. If hypotheses are not exchangeable, this may result in reduced interpretability. Consider for example the situation that $M, N \subset \{1, \ldots, m\}$ are index sets of hypotheses with distinct interpretations, e.g. hypotheses that are a priori more and less likely to be true. A user applying an FDR-controlling method will find a rejected set B, but will explicitly or implicitly split this for interpretational purposes in $B \cap N$ and $B \cap M$. FDR control, however, is not simultaneously guaranteed for these two sets, and may well be violated precisely for the smaller, and therefore often more interesting set (Finner and Roters,

2001; Goeman and Solari, 2014). In contrast, the FDP confidence bounds $q_\alpha(S \cap M)$ and $q_\alpha(S \cap N)$ are always simultaneously valid. FDP confidence bounds extend inference based on false discovery proportions to non-exchangeable settings, regardless of the complexity of the pattern of non-exchangeability.

All-resolution inference In many high-dimensional data types (e.g. neuroimaging; genomics), primary interest is not always at the level of measurement (e.g. voxels; markers), but frequently centers on higher aggregated 'regions' (e.g. brain regions; genes or pathways). Hypothesis tests for these aggregated regions are usually formulated as tests for intersection hypotheses (Goeman and Bühlmann, 2007; Nichols, 2012). The closed testing procedure joins together all regions of all sizes, including singletons, into a single multiple testing procedure in a natural way. A user interested in regions can report an adjusted p-value $\tilde{p}_S$ for each region S or, perhaps more interestingly, the confidence bound $1 - q_\alpha(S)$ as a measure of the percentage of the region S that is confidently perturbed. As these bounds are simultaneous over all regions, the user does not have to specify regions of interest in advance, but may select those freely after seeing the data without compromising the validity of the inference. Moreover, no further multiple testing correction is needed if multiple regions are selected. The concentration set Z may serve as guidance for the selection of regions.

The algorithms described in this paper have been implemented in the R-package `cherry`, available on CRAN. An illustration of the application of the methods, based on earlier versions of the algorithms, can be found in Goeman and Solari (2014).

References

Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological) 57*(1), 289–300.

Benjamini, Y. and D. Yekutieli (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics 29*(4), 1165–1188.

Bittman, R. M., J. P. Romano, C. Vallarino, and M. Wolf (2009). Optimal testing of multiple hypotheses with common effect direction. *Biometrika 96*(2), 399–410.

Brannath, W. and F. Bretz (2010). Shortcuts for locally consonant closed test procedures. *Journal of the American Statistical Association 105*(490), 660–669.

Chi, Z. (2007). On the performance of FDR control: constraints and a partial solution. *The Annals of Statistics 35*(4), 1409–1431.

Dunn, O. J. (1961). Multiple comparisons among means. *Journal of the American Statistical Association 56*(293), 52–64.

Efron, B. (2012). *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*. Cambridge: Cambridge University Press.

Finner, H. and M. Roters (2001). On the false discovery rate and expected type I errors. *Biometrical Journal 43*(8), 985–1005.

Finner, H., M. Roters, and K. Strassburger (2014). On the Simes test under dependence. *Statistical Papers*, 1–15.

Fredman, M. and M. Saks (1989). The cell probe complexity of dynamic data structures. In *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing*, STOC '89, New York, NY, USA, pp. 345–354. ACM.

Gabow, H. and R. Tarjan (1985). A linear-time algorithm for a special case of disjoint union. *Journal of Computer and System Sciences 30*(2), 209–221.

Genovese, C. R. and L. Wasserman (2006). Exceedance control of the false discovery proportion. *Journal of the American Statistical Association 101*(476), 1408–1417.

Goeman, J. J. and P. Bühlmann (2007). Analyzing gene expression data in terms of gene sets: methodological issues. *Bioinformatics 23*(8), 980–987.

Goeman, J. J. and A. Solari (2011). Multiple Testing for Exploratory Research. *Statistical Science 26*(4), 584–597.

Goeman, J. J. and A. Solari (2013). Correction note. *Statistical Science 28*(3), 464–464.

Goeman, J. J. and A. Solari (2014). Multiple hypothesis testing in genomics. *Statistics in Medicine 33*(11), 1946–1978.

Gou, J., A. C. Tamhane, D. Xi, and D. Rom (2014). A class of improved hybrid Hochberg–Hommel type step-up multiple test procedures. *Biometrika 101*(4), 899–911.

Henning, K. S. and P. H. Westfall (2015). Closed testing in pharmaceutical research: Historical and recent developments. *Statistics in Biopharmaceutical Research 7*(2), 126–147.

Hochberg, Y. (1988). A sharper Bonferroni procedure for multiple tests of significance. *Biometrika 75*(4), 800–802.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics 6*(2), 65–70.

Hommel, G. (1986). Multiple test procedures for arbitrary dependence structures. *Metrika 33*(), 321–336.

Hommel, G. (1988). A stagewise rejective multiple test procedure based on a modified Bonferroni test. *Biometrika 75*(2), 383–386.

Hommel, G. (1989). A comparison of two modified Bonferroni procedures. *Biometrika 76*(3), 624–625.

Lehmann, E. and J. Romano (2005). Generalizations of the familywise error rate. *The Annals of Statistics 33*(3), 1138–1154.

Marcus, R., E. Peritz, and K. Gabriel (1976). Closed testing procedures with special reference to ordered analysis of variance. *Biometrika 63*(3), 655–660.

Nichols, T. E. (2012). Multiple testing corrections, nonparametric methods, and random field theory. *Neuroimage 62*(2), 811–815.

Pacifico, M. P., C. Genovese, I. Verdinelli, and L. Wasserman (2007). Scan clustering: a false discovery approach. *Journal of Multivariate Analysis 98*(7), 1441–1469.

Rødland, E. A. (2006). Simes' procedure is 'valid on average'. *Biometrika 93*(3), 742–746.

Sarkar, S. (2008). On the Simes inequality and its generalization. *IMS Collections Beyond Parametrics in Interdisciplinary Research: Festschrift in Honor of Professor Pranab K. Sen 1*, 231–242.

Simes, R. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika 73*(3), 751–754.

Tarjan, R. (1975). Efficiency of a good but not linear set union algorithm. *Journal of the ACM 22*(2), 215–225.

Wright, S. P. (1992). Adjusted p-values for simultaneous inference. *Biometrics 48*(4), pp. 1005–1013.